SEVENTH SEMI-ANNUAL REPORT

ON

RESEARCH ON

CONTROL OF FREE-FLYING ROBOT MANIPULATOR SYSTEMS

Submitted to

m, Jr., Chief, Information Sciences Division
search Center, Moffett Field, CA 94035

by

University Aerospace Robotics Laboratory
ment of Aeronautics and Astronautics
ford University, Stanford, CA 94305

rformed Under NASA Contract NCC 2-333
period March 1990 through August 1991

Professor Robert H. Cannon Jr.
Principal Investigator

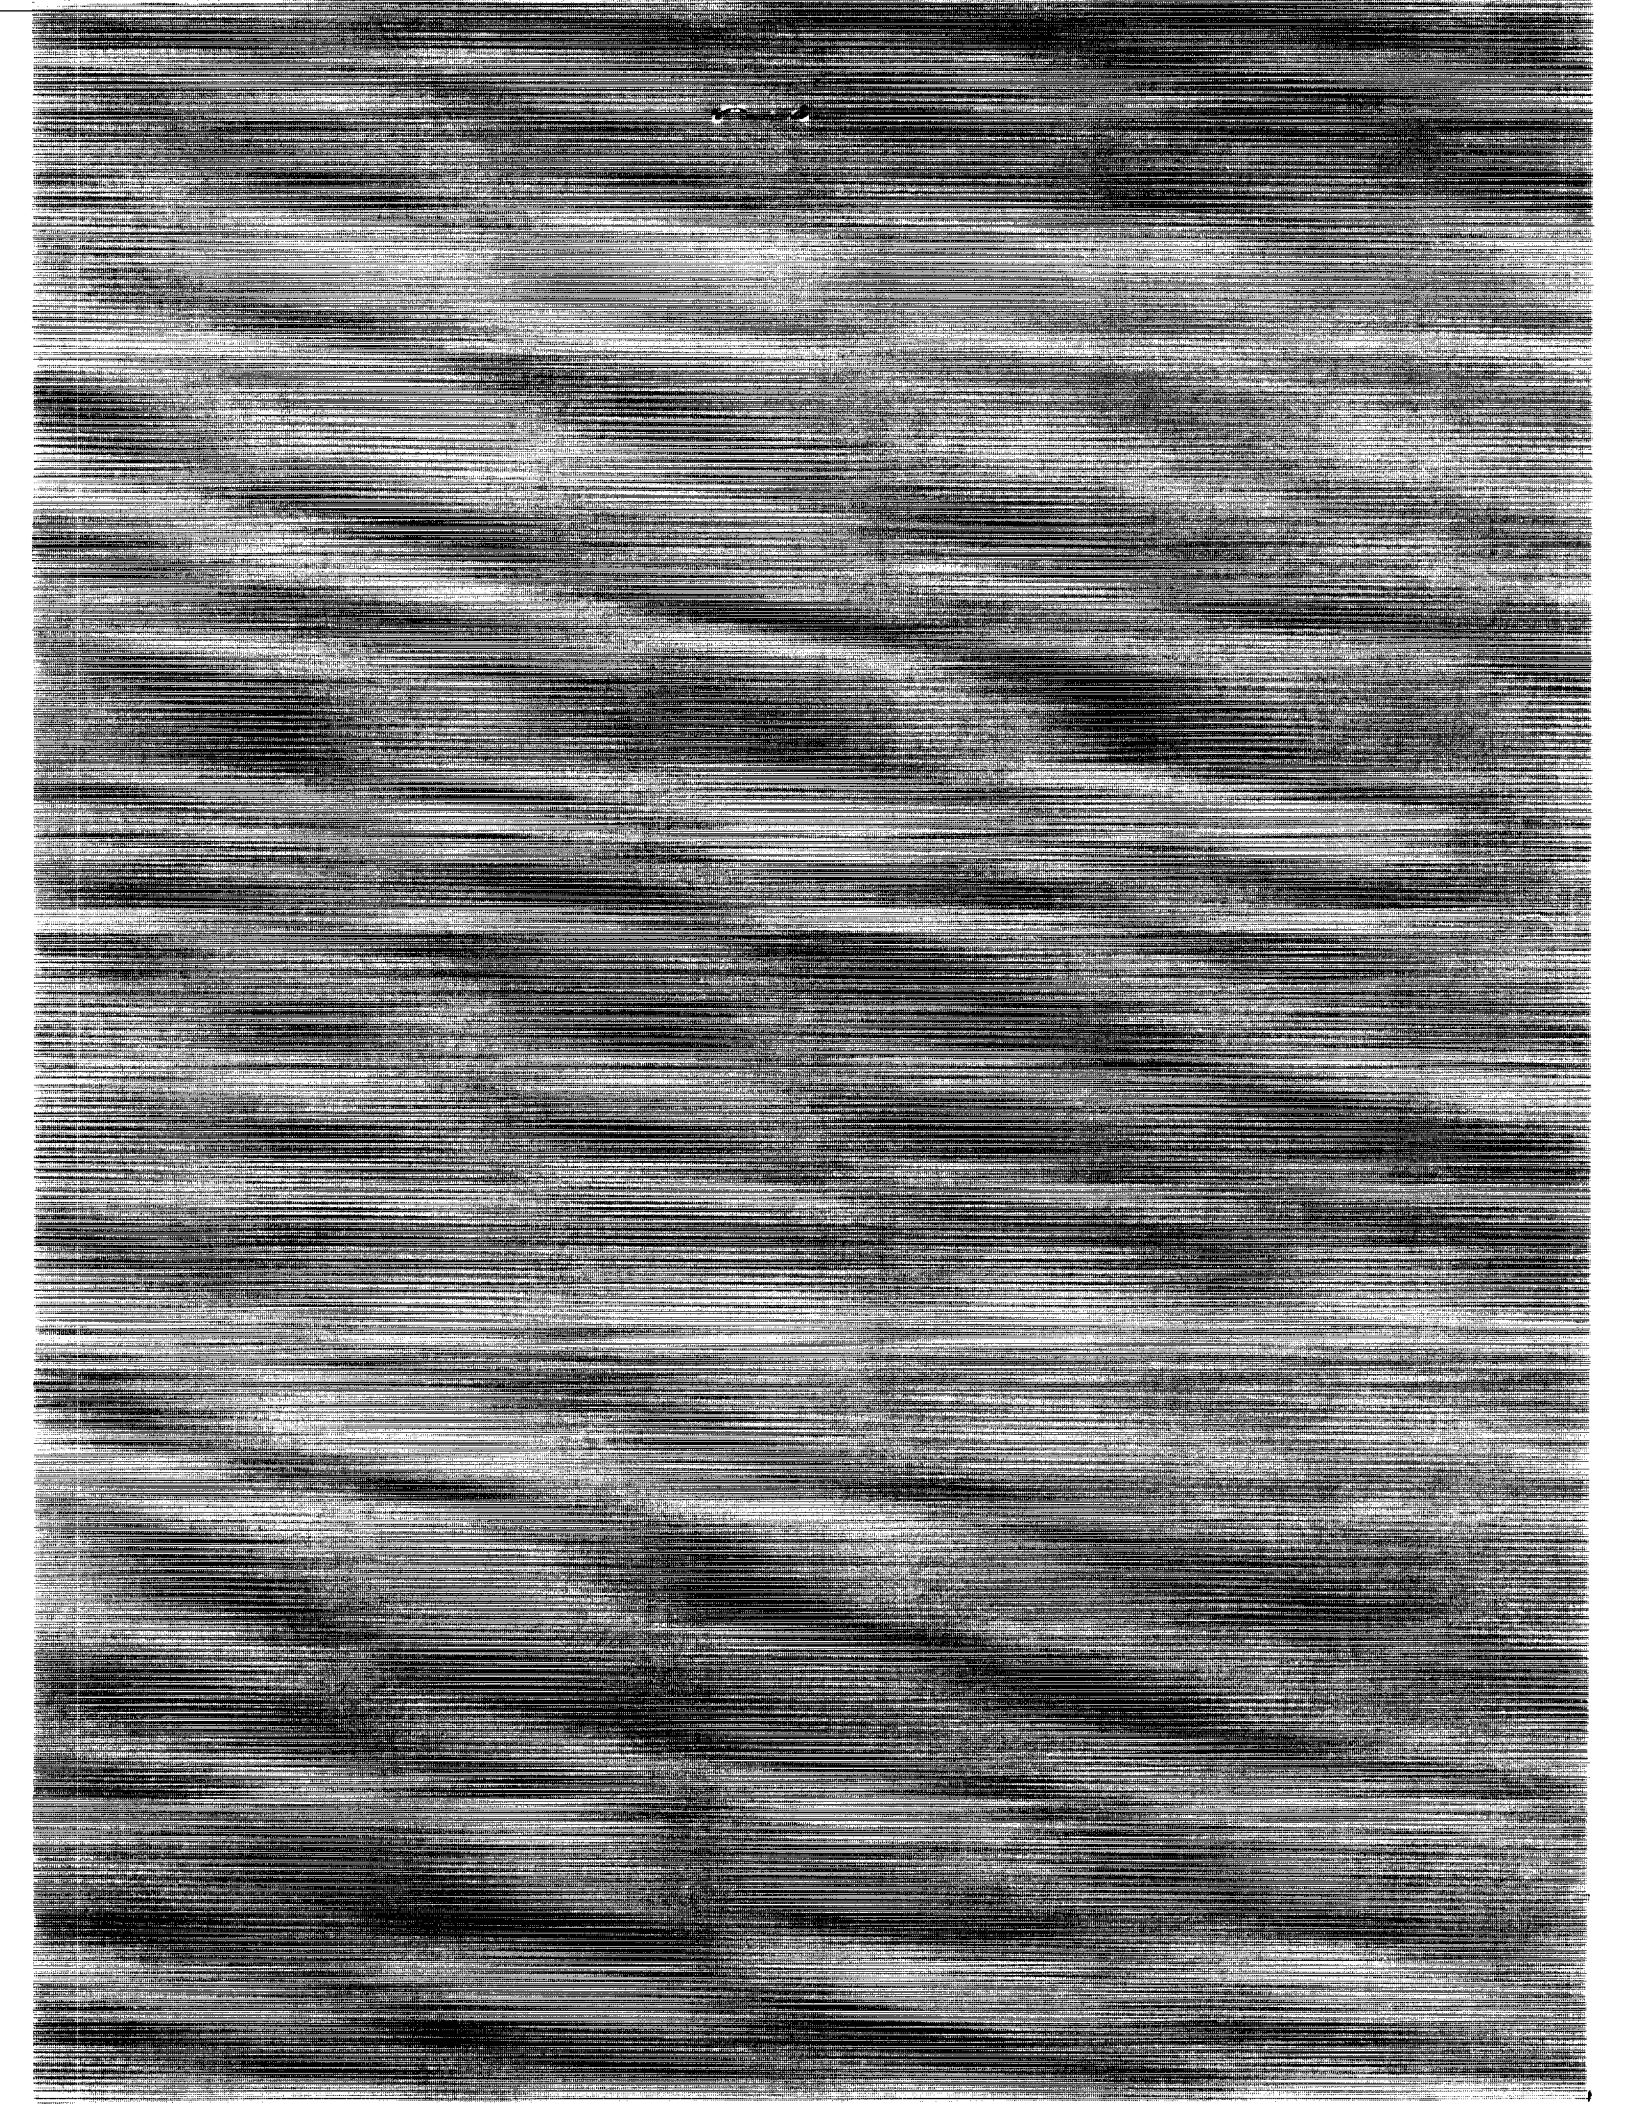

ELEVENTH SEMI-ANNUAL REPORT

ON

RESEARCH ON

CONTROL OF FREE-FLYING SPACE ROBOT MANIPULATOR SYSTEMS

Submitted to

Dr. Henry Lum, Jr., Chief, Information Sciences Division
Ames Research Center, Moffett Field, CA 94035

by

The Stanford University Aerospace Robotics Laboratory
Department of Aeronautics and Astronautics
Stanford University, Stanford, CA 94305

Research Performed Under NASA Contract NCC 2-333
During the period March 1990 through August 1991

Professor Robert H. Cannon Jr.
Principal Investigator

~

Contents

List of Figures v

1 Introduction 1

1.1 Summary of Progress 5

2 Multiple Robot Cooperation 7

2.1 Introduction 7

2.2 Progress Summary 7

2.3 Experimental Hardware 8

2.4 Control 9

2.5 Experimental Results 9

2.6 Future Work 12

3 Neural Networks for Control of Space Robots 13

3.1 Introduction 13

3.2 Planned Experimental Investigations 14

4 Multi-Robot Space Assembly 17

4.1 Introduction 17

4.2 Assembly Using Potential Field Navigation Techniques 17

4.3 Simulation Results 18

4.4 Future Work 18

5 Experiments in Control of a Flexible Arm Manipulating a Dynamic Payload 19

5.1 Introduction 19

5.2 Modelling and Coupling Analysis 20

5.3 Stability of LQR Design 23

5.4 Experimental Results 25

5.5 Conclusions and Future Work 29

Bibliography 31

A Experiments in Autonomonous Navigation and Control of a Mulit-Manipulator Free-Flying Space Robot 33

List of Figures

2.1 Two-Robot Object Manipulation 10
2.2 Object Slew Performance 11

5.1 Experimental Apparatus Schematic 21
5.2 Frequencey Response 22
5.3 Nominal Arm Mode Shapes and Frequencies 22
5.4 Straw Man Problem 23
5.5 LQR design w/o modelling pendulum 24
5.6 LQG design w/o modelling pendulum 24
5.7 Free Motion of Pendulum 26
5.8 Response With Dynamic Payload 26
5.9 Response With Dynamic Payload 27
5.10 Response With Dynamic Payload 28

~

Chapter 1

Introduction

This document reports on the work done under NASA Cooperative Agreement NCC 2-333 during the period March 1990 through August 1991 . The research was carried out by a team of five Ph.D. candidate students from the Stanford University Aerospace Robotics Laboratory under the direction of Professor Robert H. Cannon, Jr. The goal of this research is to develop and test experimentally new control techniques for self-contained, autonomous free-flying space robots. Free-flying space robots are envisioned as a key element of any successful long term presence in space. These robots must be capable of performing the assembly, maintenance, inspection, and repair tasks that currently require astronaut extra-vehicular activity (EVA). Use of robots will provide economic savings as well as improved astronaut safety by reducing and in many cases eliminating the need for human EVA.

The focus of our work is to develop and carry out a set of research projects using laboratory models of satellite robots and a flexible manipulator. The second-generation space-robot models use air-cushion-vehicle (ACV) technology to simulate in two dimensions the drag-free, zero-g conditions of space. Using two large granite surface plates (6' by 12' and 9' by 12') which serve as the platforms for these experiments, we are able to reduce gravity-induced accelerations to under $10^{-5}g$, with a corresponding drag-to-weight ratio of about 10^{-4}—a very good approximation to the actual conditions in space. The flexible manipulator, also using air-cushion technology, is mounted on a third (4' by 8') granite surface plate.

The Global Navigation and Control project demonstrates simultaneous control of the robot manipulators and the robot base position on the free-flying robot model. This allows manipulation tasks to be accomplished while the robot body is controlled along a trajectory. The robots have the capability to rendezvous and capture free-flying objects. Once the objects have been captured, the robots can manipulate and/or transport them to their desired location. The robot actions are directed by a user who uses a graphical interface to issue high-level commands such as "capture" to the robot. This project has been completed and is in the documentation phase.

Having already demonstrated global navigation and control of a free-floating robot, current work is divided into four major research projects: Multiple-Robot Cooperation, Neural

Networks for Control of Space Robots, Navigation for Multiple Robots Doing Space Assembly, and Dynamic Payload Manipulation. Each of these projects represents an ongoing experimental PhD thesis.

The Multiple-Robot Cooperation project will demonstrate multiple free-floating robots working in teams to carry out tasks too difficult or complex for a single robot to perform. A third space robot model, identical to the robot fabricated for the Thrusterless Locomotion project, recently has become operational– providing the minimal two robots needed for the multiple-robot research.

The goal of the Neural Networks project is to do the same task as the Global Navigation and Control project with the use of Neural Networks instead of model based control techniques. The Neural Networks will imbue the robots with the capability to perform well in the presence of effects such as bias torques and friction which are difficult for model based controllers to handle. The Neural Networks should also allow the robot handle unstructured environments better than the existing controllers.

The Navigation for Multiple Robots Doing Space Assembly project is concerned with coordinating the moves of many free flying robots assembling a space truss. The idea here is to use a potential field approach to prevent the robots from colliding with each other or other obstacles. The potential field acts like a air traffic controller directing all of the robots to do an orderly and efficient assembly.

The Dynamic Payload Manipulation project seeks to demonstrate control of non-rigid payloads and explore the payload's effects on the dynamics of a manipulator system. This research addresses the fundamental issues involved with manipulating space-born objects that possess sloshing fuel tanks or flexible appendages such as solar arrays.

The chapters that follow give detailed progress and status reports on a project-by-project basis.

Included with this report are six recently published papers.

1. Warren J. Japer and Robert H. Cannon, Jr., "Initial Experiments in Thrusterless Locomotion Control of a Free-Flying Robot, " *Proceedings of the ASME Winter Annual Meeting,* Dallas, Texas, November 1990.

2. Harold L. Alexander and Robert H. Cannon, Jr., "An Extended Operational-Space Control Algorithm for Satellite Manipulators," *The Journal of the Astronautical Sciences,* Vol. 38, No. 4, October-December 1990, pp 473-486.

3. Stanley A. Schneider, Marc A. Ullman, Vincent W. Chen, "ControlShell: A Real-Time Software Framework," *Proceedings of the IEEE Conference on Systems Engineering,* Dayton, Ohio, 1991.

4. Wen-Wei Chiang, Raymond Kraft and Robert H. Cannon, Jr., "Design and Experimental Demonstration of Rapid, Precise End-Point Control of a Wrist Carried by a Very Flexible Manipulator," *The International Journal of Robotics Research,* Vol. 10, No. 1, February 1991.

5. R. Koningstein and R. H. Cannon, Jr., "Experiments with Simplified Computed-Torque Controllers for Free-Flying Robots,," *Proceedings of the American Control Conference,* Boston, Massachusetts, June 1991.

6. L. J. Alder and S. M. Rock, "Control of a Flexible Robotic Manipulator with Unknown Payload Dynamics: Initial Experiments," to be presented at the *ASME Winter Annual Meeting,* Atlanta, Georgia, November 1991.

~

1.1 Summary of Progress

- Advanced the capability for cooperative manipulation with multiple robots under the management of a coordinating agent.
- Demonstrated ability for multiple robots to rendezvous and capture a single large free-floating object.
- Showed that if internal dynamics of a payload are ignored, conventional control techniques may cause instability.
- Demonstrated ability of a flexible robot arm to control an object with internal dynamics without direct measurement of the internal dynamics. This includes the ability to damp out the internal dynamics of the object.
- Defined applications for Neural Networks in the control of space robots.
- Simulated the use of potential field navigation techniques to control the motions of many robots doing a space assembly.

~

Chapter 2

Multiple Robot Cooperation

William C. Dickson

2.1 Introduction

This chapter summarizes our progress to date in the area of multiple robot cooperation. This work will eventually unite the various lines of research presently being conducted in fixed- and floating-base cooperative manipulation, and in global navigation and control of space robots. Our goal is to demonstrate multiple free-floating robots working in teams to carry out tasks too difficult or complex for a single robot to perform. Achieving this cooperative ability will involve solving specialized problems in dynamics and control, high-level path planning, and communication.

2.1.1 Research Goals

Some of the goals of this project are:

- Cooperative object capture and manipulation by a robot team.
- Fine cooperative manipulation in presence of on-off control.
- Path generation considering dynamic constraints and obstacle avoidance.
- Development of control strategies for path following.

2.2 Progress Summary

Activities completed from March 1991 to August 1991 were:

- Incorporated use of Condition Module for team synchronization.
- Added object capture to capabilities of robot team.
- Force sensors mounted on robot grippers.

- Off-board battery charger designed and built.

2.3 Experimental Hardware

2.3.1 Overview

The experimental hardware associated with this research currently consists of an off-board vision system, two mobile robots, an off-board coordinator processor, and a manipulation object. The robots and object use self-powered air bearings for flotation on a 6' by 12' granite table. An off-board battery charger/discharger was recently developed.

2.3.2 Vision System

The vision system consists of a camera mounted above the granite table, an ARL-developed Point Grabber Vision board [2], and a commercial 68030-based computer for vision processing. The vision board converts camera bright spots into a list of pixel coordinates and intensity values. The bright spots are produced by infrared (IR) light emitting diodes (LEDs) located on the robots and object. The vision computer uses the information generated by the point-grabber board to determine the positions of the robots, the robot manipulator endpoints, and the object(s) [3].

2.3.3 Robots

The robots used in this research are nearly identical to the original second-generation robot currently used in the Navigation and Control research. One major differences is that these robots utilize a momentum wheel — allowing the robots to control their orientation without the use of thrusters. Second, unlike the original second-generation robot, these robots currently have no on-board vision system for improved workspace sensing.

The grippers are pneumatically driven plungers used by the robots to manipulate floating objects. These grippers feature commercial linear bearings for precise, repeatable performance. Strain-gage force sensors were recently installed on the grippers. The strain gages are mounted in opposing pairs for each direction of sensing to attain differential signals (to reject temperature effects). The gripper on each arm is equiped to measure forces in two perpendicular horizontal directions. The force sensors will be used to close the control loop that supplies forces for object manipulation and to help identify arm-joint torque biases due to wires, tubes, and joint friction.

2.3.4 Manipulation Object

The manipulation object is constructed from two half-square-feet floating pads connected by a three-feet-long metal bar. Four cylindrical grip ports on the object facilitate grasping by the robots. Battery powered IR LEDs allow the object to be tracked by the vision system.

2.3.5 Off-board battery charging

An off-board battery charger/discharger unit has been constructed. The charging is current-rather than voltage-controlled as required by the nickle-cadmium batteries. The device is also used to discharge malfunctioning batteries at a controlled current. The off-board charging and discharging capability provides extended robot operation time, since the robots' on-board charging occurs only when the robots are immobilized by a removable external power cable.

2.4 Control

2.4.1 Task synchronization

The Condition Module discussed in the Tenth semi-annual report [1] has recently been incorporated into the control scheme of the robot team. This module allows the robots and coordinator to transmit and receive high-level control messages to and from one other. Using these messages, the robots are able to determine when to approach, grasp, move, and release the object.

2.5 Experimental Results

Previous experiments have successfully demonstrated the robot team manipulating a floating object. In these experiments, an off-board vision system tracks the positions and velocities of the object and two robots (as well as the robots' manipulator endpoints), and sends this information via the network to each of the robots. The robots move the object to a location specified by the operator via a user interface. The on-board air thrusters and momentum wheels control each robot's position and orientation as described. Figures 2.1 and 2.2 show the results of an example object slew. The translational errors were regulated to sub-millimeter levels, while the rotational error was less than one degree. Note that the step changes in the "desired" positions indicate the destination state of the object. As discussed in [1], the object is regulated to a position near the robots when the robots are out of range of the destination state. Once in range, the robots regulate themselves and the object to the destination state. In the example slew, the robots came into range at the 18 second point of the run, or about 14 seconds after the new destination state was commanded. These experiments began with the object already in the grasp of the team.

The incorporation of the Condition Module has provided the additional capability of a robot-team object capture. With this capability, the user merely chooses the object for capture via the graphical user interface, and the robot team moves into position and makes the grasp. The user can then specify where the object should be placed, and the team carries out the move autonomously. After placing the object, the robots are ready for the next user command.

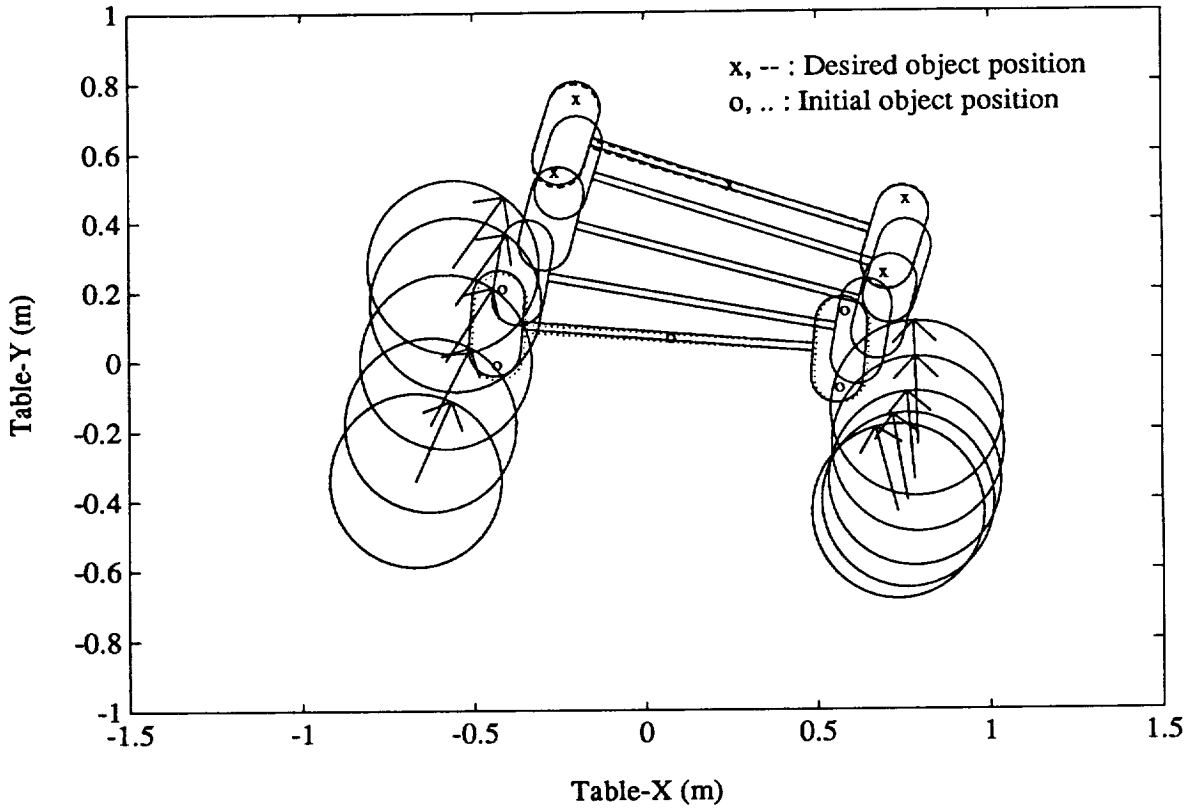

Figure 2.1: Two-Robot Object Manipulation

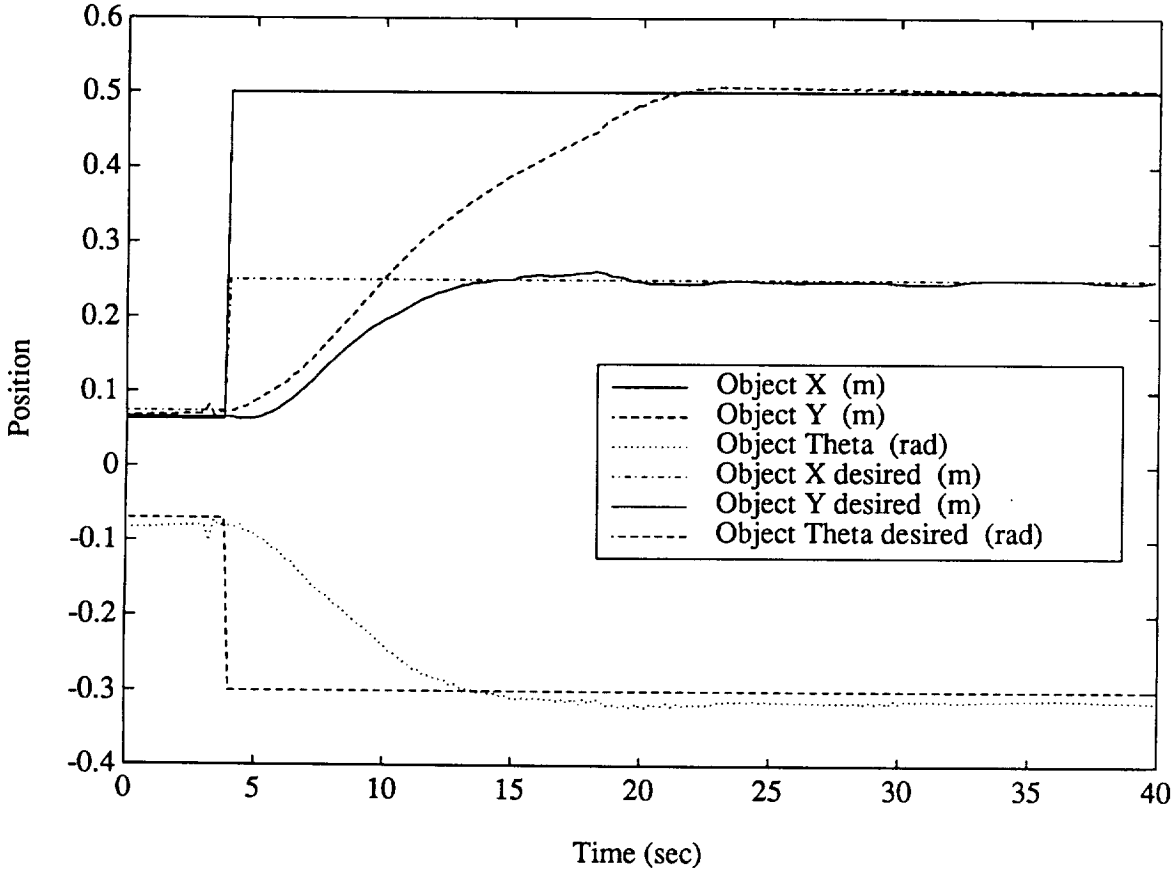

Figure 2.2: Object Slew Performance

2.6 Future Work

Although the control methods used in this research have proven successful in experiments, a more detailed theoretical analysis is required to fully understand the generic issues and how the methods can be applied to other systems.

The addition of the force sensors will allow us to actually measure the forces being exerted on the object by the grippers. Presently, the control inputs are assumed to provide these forces exactly. With the measurements, the controls can be modified to counteract unmodelled spring and friction torques at the joints. These steps would directly improve the robot team's ability to precisely control the grasped object.

The following hardware issues remain:

- Calibrate the new force sensors.
- Attain wireless Ethernet communication to replace fiber-optic and coaxial cables.

Chapter 3

Neural Networks for Control of Space Robots

Edward Wilson

3.1 Introduction

Because they are capable of complex learned behaviors, adaptive neural networks have the potential to make a significant impact on the field of robotics in the near future. To investigate this potential, the ARL is launching a new program of experimental research to examine the applicability of this exciting technology to the control of space robot manipulator systems.

Neural networks are loosely modeled after the human brain. Instead of performing calculations sequentially on a single processor, calculations are performed simultaneously (even asynchronously) by a network of relatively simple processors. These processors act only locally, producing a single output based on a limited number of inputs (often the outputs of neighboring processors), just as the neurons in a human brain do.

Networks of these simple processors have emergent properties that allow very complex behavior–such as learning and pattern recognition–that are presently very limited in current computers. For example, neural networks may be used to implement arbitrary mappings of inputs to outputs, such as sensor signals to actuator commands. Since the mapping can be taught indirectly, neural networks are especially attractive for poorly-understood systems; they can generalize from training inputs and then respond in untaught situations. Due to the distributed nature of the processing, networks are often robust to internal component failures; the remaining processing elements can adapt to account for the failure. Similarly, the network can be made to adapt to changes in the environment, plant, performance criteria, etc.

One significant advantage of neural networks is that they may ultimately be implemented on parallel processing hardware for greatly enhanced throughput capabilities; however, they are often implemented on traditional sequential processing computers during

development, and when processing speed is not a limiting factor.

3.2 Planned Experimental Investigations

Our initial experiments will employ this strategy of using existing "sequential" computers. In particular, existing Motorola 68030 processors will be used, with possible upgrades to Motorola 68040, DSP, or dedicated neural processing hardware - as computation requirements are determined by our initial investigations.

In "supervisory learning", the parameters in a network are chosen by training it to emulate another controller. In a series of applications we could employ supervisory learning in training a network to emulate proportional-integral-derivative, bang-bang, computed torque, or even human controllers on an existing space robot manipulator system. This relatively simple training technique will yield important information about learning capabilities and the computational requirements for more sophisticated neural controllers. Ultimately, we plan to use non-supervisory learning, in which instead of adapting the network to mimic an existing controller it attempts to please a critic (optimizing a grading score, for example).

The training procedures will first be developed in simulation, then controllers will be trained in simulation until steady state is reached. The final training will be performed on the actual plant to take into account the existing non-linearities and modeling errors.

3.2.1 Initial Experiment: System Bias Identification and Correction

As a first step, a neural network could augment an existing arm controller by calculating bias torques (due to friction, motor bias, wires, hoses, etc.).

A neural network is well-suited to perform this bias torque mapping because:

- The problem requires some form of learning.
- Sources of the bias are not fully understood.
- The bias mapping is sure to be non-linear.
- The mapping could be time-varying (especially if the hoses and wires shift around) which would require some sort of on-line adaptation.

As well as being a useful augmentation to existing and future controllers, this relatively straight forward preliminary task will yield important information about the efficiencies of various network structures and training algorithms, as well as the computational requirements for more sophisticated neural controllers.

3.2.2 Possible Further Experiments: Capture of a Free-Floating Object

Experiments will be performed on the space robot named Heavenly. This robot has recently accomplished tracking, capture, and manipulation of a free-floating object, as discussed in the chapter by Marc Ullman.

As a goal that will demonstrate the capabilities of the neural controller to be developed, we hope to achieve performance that is similar to the existing controller, yet enhanced both due to the intrinsic adaptive nature of the neural networks and the obstacle avoidance that the proposed neural navigation scheme will provide.

Following is a proposed plan of attack:

1. Base position and attitude control using thrusters:

   First, a regulator, then a potential field navigator will be developed, hopefully through non-supervisory training - perhaps using back-propagation through time as has been developed by Prof. Widrow's group at Stanford. This appears to be a good first step since the two translational and one rotational degrees of freedom can be decoupled (after pre-processing to make the thrusters appear linear) into simple double-integrator plants.

2. Target intercept:

   First a fixed, then a moving target will be tracked. If a potential field navigator is used, obstacles could be avoided by simply augmenting the potential field.

3. Target grasp:

   To simplify the problem, one arm will be used, and a single grip point will be mounted at the mass center of the object. The arm endpoint to target grasp point intercept is a similar concept to the base to target c.g. intercept, yet it is much more difficult due to the nonlinear arm dynamics as well as the problem of multiple inputs and outputs. Using non-supervisory learning to accomplish this MIMO control task is a very challenging problem.

Chapter 4

Multi-Robot Space Assembly

Kurt R. Zimmerman

4.1 Introduction

The inherent danger of the space environment make it more attractive to use robots rather than humans for a significant number of tasks. Future utilization of space will require the ability to assemble modular payloads which are too large to be launched in one piece (for example, the space station) and will require many hours of high risk EVA. Sophisticated robots which are more suited to the environment may prove to be the key factor in determining the success of such assembly tasks.

The Stanford ARL has already investigated several of the major aspects of robotic space assembly, the most relevent being the work of Ullman and Dickson in [2]. Also particularly relevent is the recent work of [5] and [4]. A demonstration of space assembly could be achieved by incorporating the results of these research efforts into an appropriate multi-robot control regime.

4.2 Assembly Using Potential Field Navigation Techniques

The initial thrust of this research was to identify a suitable technique for controlling several mobile robots operating within the same workspace. A potential field navigation technique was chosen since it requires little planning overhead and is robust in the face of unexpected events. The concept of potential field navigation is that the robot creates a potential field map of the workspace in which all obstacles appear as large potential barriers while the robot's goal appears as an attractive potential well. The robot can then navigate through the obstacles in its environment by simply following the gradient of the potential field down to the goal. The goal may be an object that the robot is fetching or the desired destination of an object that the robot is carrying. Barriers in the potential field may be structural objects or other robots (therefore, the barrier locations can be dynamically changing).

Conventionally, the potential field is created by summing an attractive and a repulsive

field. The attractive field is based on either a squared or linear distance to goal (a parabolic or conical well) and the repulsive field is based on the inverse squared distance to the obstacle. A common problem that arises with this technique is that the potential field often contains local minima, causing the robot to navigate to an incorrect location and get stuck. A technique recently developed at Stanford (by Edward S. Plumer, a current student in the Electrical Engineering PhD program) will generate a potential field which is guaranteed to have no local minima for the case of a single point robot with a single goal. In this technique the potential field is created by first creating a barrier field which decays geometrically away from obstacles, and then cummulatively summing the distance to the goal with the barrier field to achieve the final potential field. This was the technique used in the multiple robot simulations described below.

4.3 Simulation Results

To test the effectiveness of the potential field navigation technique, a simulator was designed in which an arbitrary number of robots could be placed in a given workspace along with an arbitrary number of floating objects (structural pieces). The simulation operator could designate a desired complete structure made from the floating pieces. The robots would then cooperatively retrieve the objects and build the structure.

Some of the advantages of the potential field technique were found to be that it requires very little planning on the part of each individual robot and it is quite robust with respect to the dynamic interactions of the multiple robots. However, the computation overhead for the potential field is quite high and there are instances in which the robots can get stuck in minima caused by the presence of other robots (the guarantee of no local minima is only for a single robot in the workspace with a single goal).

4.4 Future Work

The near term goals for this research are to improve the potential field technique using the simulator. Methods for incorporating a high-level manager for handling situations in which the potential technique fails will also be investigated. The long term goal of this research is to implement the multi-robot controller on the existing hardware in the lab to demonstrate a space assembly task. Several modifications to the existing hardware will have to be made. Most critical to the success of this task will be the implementation of a potential field generator processor/multi-processor which is capable of generating the potential fields in real time.

Chapter 5

Experiments in Control of a Flexible Arm Manipulating a Dynamic Payload

Lawrence J. Alder

5.1 Introduction

For many applications, flexible robot arms may handle payloads that are not simple rigid bodies. In space applications, the RMS (remote manipulator system) will be manipulating satellites that may contain fuel or have flexible appendages. Most high performance control schemes for flexible manipulators require end point feedback. Such control systems have been shown to be sensitive to unmodelled dynamics in the payload. If the payload dynamics are not accounted for in the control design, degraded performance and instability are possible.

An experimental apparatus has been constructed to investigate the problem of controlling a dynamic payload with a flexible manipulator robot. The robot is a single flexible link that operates in a horizontal plane. The dynamic payload is a box with a pendulum inside of it. The apparatus has been designed to encompass the fundamental issues of the problem without undo complexity.

This chapter details the research goals and the experimental progress in controlling the above mentioned system.

5.1.1 Research Goals

The goals of this project are[1]:

[1]These goals are a reprint from the NASA Ninth Semi-Annual Report.

- Demonstrate stable end-point control of a flexible arm manipulating both rigid and nonrigid objects
- Demonstrate the ability of a flexible robot arm to damp internal vibrations in a payload
- Demonstrate the ability to move a nonrigid object from an initial point to a final point with no residual vibrations
- Demonstrate all of the above without knowledge of whether the payload is rigid

5.1.2 Research Progress

- Have demonstrated stable end-point control of a flexible arm manipulating both rigid and nonrigid objects
- Have demonstrated the ability of a flexible robot arm to damp internal vibrations in a payload
- Have demonstrated the ability to move a nonrigid object from an initial point to a final point with no residual vibrations
- Have accurately modelled the experimental system including effect of dynamic payload
- Have shown that if the payload dynamics are ignored it is likely that linear quadratic regulator control will destabilize the payload dynamics

5.2 Modelling and Coupling Analysis

Figure 5.1 shows the hardware configuration. An accurate mathematical model of the hardware is constructed to aid in control design. The finite element technique is used to model the system. The accuracy of a model is validated by comparing the model with experimental data. Figure 5.2 shows a comparison of the finite element model with experimental data taken with a frequency response analyzer. Three outputs are measured: the hub angle, the pendulum angle, and the inertial endbody[2] angle. The phase plots are not shown but the agreement between experiment and model is just as good as the magnitude plots.

Using the finite element model, the mode shapes of the arm can be plotted. Figure 5.3 shows the flexible mode shapes of the arm. The rigid body shape is not shown. The shapes are scaled to the mass matrix of the system. Each plot is titled with the eigenvalue element corresponding to pendulum deflection for that frequency. The plots each are labeled with the natural frequency of that mode in Hertz. The first three plots have a solid line corresponding to a rigid payload and a dashed line corresponding to the model when the payload is dynamic. The "*" indicates where the flexible beam attaches to the endbody.

[2]Endbody is any rigid body that is rigidly attached to the endpoint of the arm.

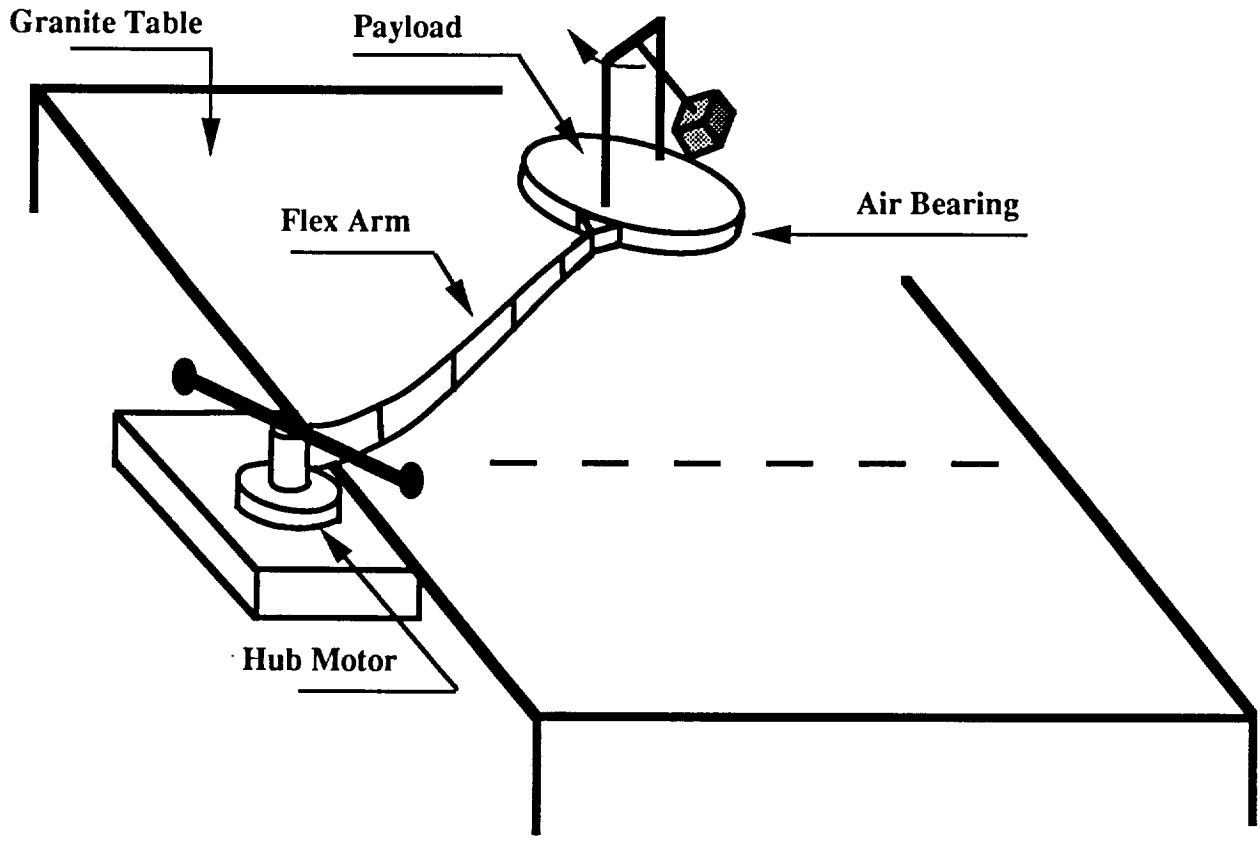

Figure 5.1: **Experimental Apparatus Schematic**

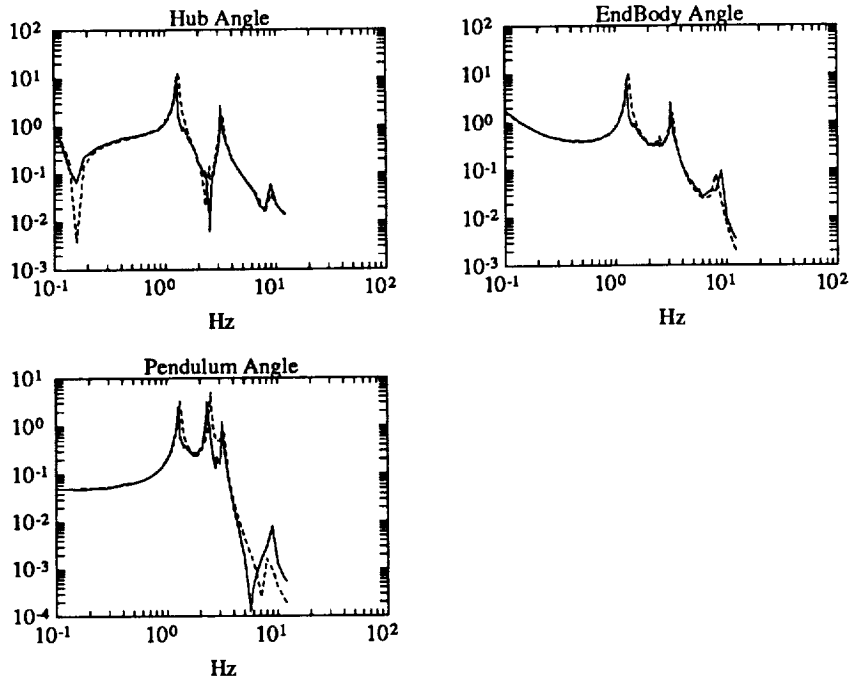

Figure 5.2: **Frequencey Response** *Solid: Exper. Dashed: Finite Elem*

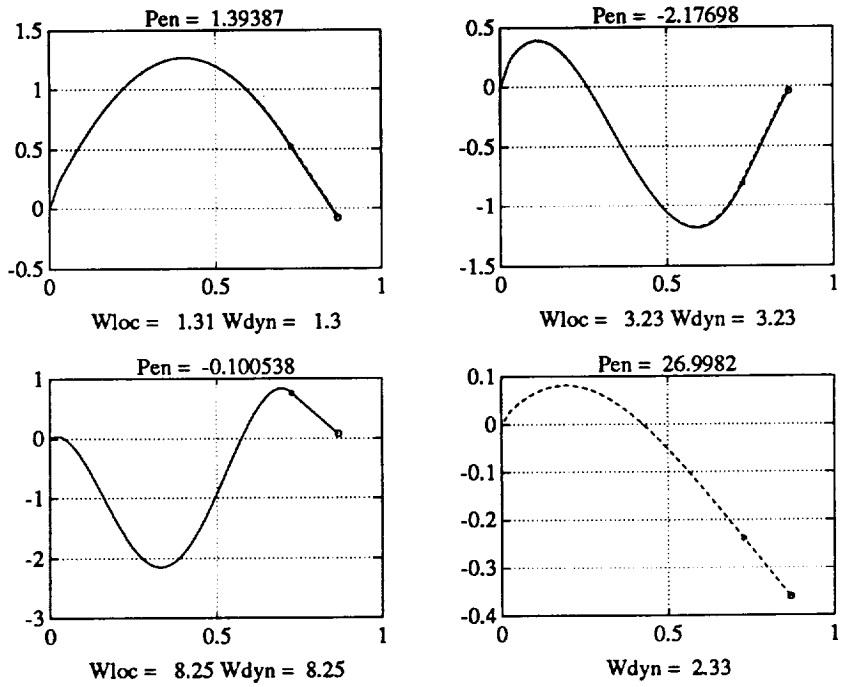

Figure 5.3: **Nominal Arm Mode Shapes and Frequencies**

solid-locked dashed-unlocked

The "o" indicates the center of mass of the endbody. The fourth plot shows the extra mode and shape introduced by the dynamics of the payload.

From Figure 5.3, one sees that the coupling of the system is primarily from the payload to the arm. In other words, the first three arm modes (1.3, 3.2 and 8.25 Hertz) and mode shapes are uneffected by the existence of the pendulum. However, the fourth plot shows that when the pendulum is excited the arm will also become excited. Again, for the first three modes the arm will be excited but not the pendulum, but for the fourth mode both the arm and pendulum become excited. This indicates that the addition of the pendulum couples into the arm dynamics but the arm dynamics do not greatly effect the pendulum dynamics.

The observability of the system from different sensors can also be examined from Figure 5.3. Imagine having a single vision sensor at the center of mass of the endbody (where the "o" is on the plots). The observability of a mode by this sensor is found by looking at the amount the center of the endbody deflects for each mode. The first beam modes (1.3, 3.2 and 8.25 Hertz) will barely be observed compared to the payload mode (2.3 Hertz). However, if we sense the inertial angle of the endbody, we will have good observability of all of the flexible modes.

5.3 Stability of LQR Design

The control problem, in its most general sense, is to design a regulator for a linear plant which has some unknown dynamics. Figure 5.4 shows the problem schematically. In Figure 5.4, Block #2 has an unknown portion of its mass on the pendulum. Hence, a control designer only knows the sum of the mass of the pendulum and second cart. He would not know the length of the pendulum or the mass ratio between the cart and pendulum.

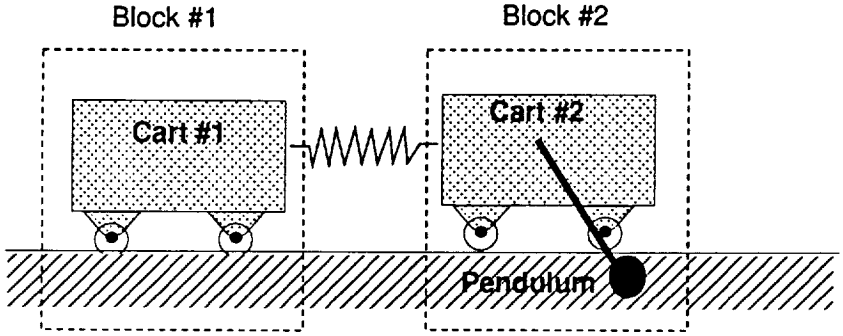

Figure 5.4: **Straw Man Problem**

It is interesting to ask what will happen when one tries to design an LQR controller for the system without knowing the pendulum exists. Figure 5.5 shows a locus of closed loop roots when a design is done for a case with the pendulum locked but the pendulum is really free to oscillate. The locus is plotted versus the length of the pendulum. When the

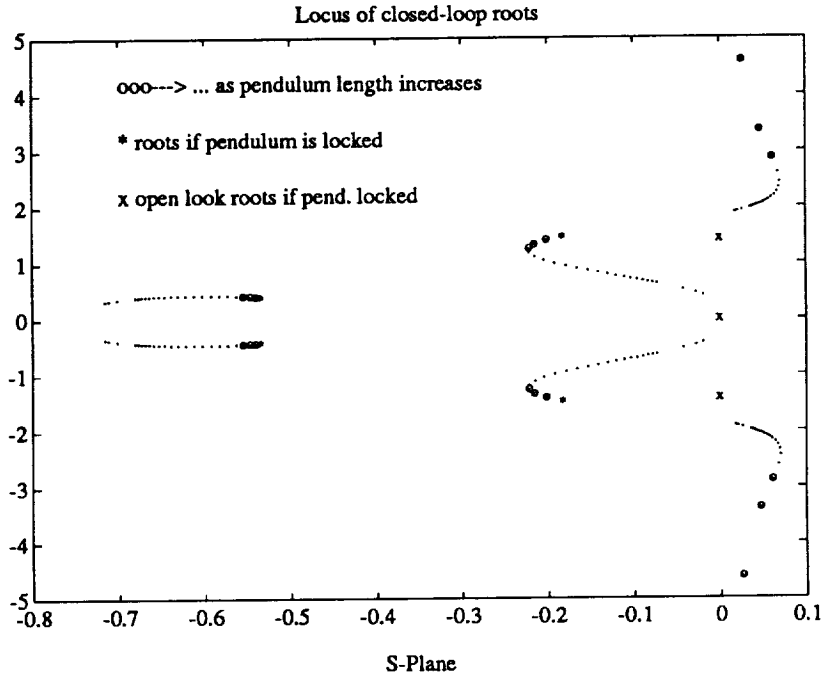

Figure 5.5: **LQR design w/o modelling pendulum** *For Straw Man Problem*

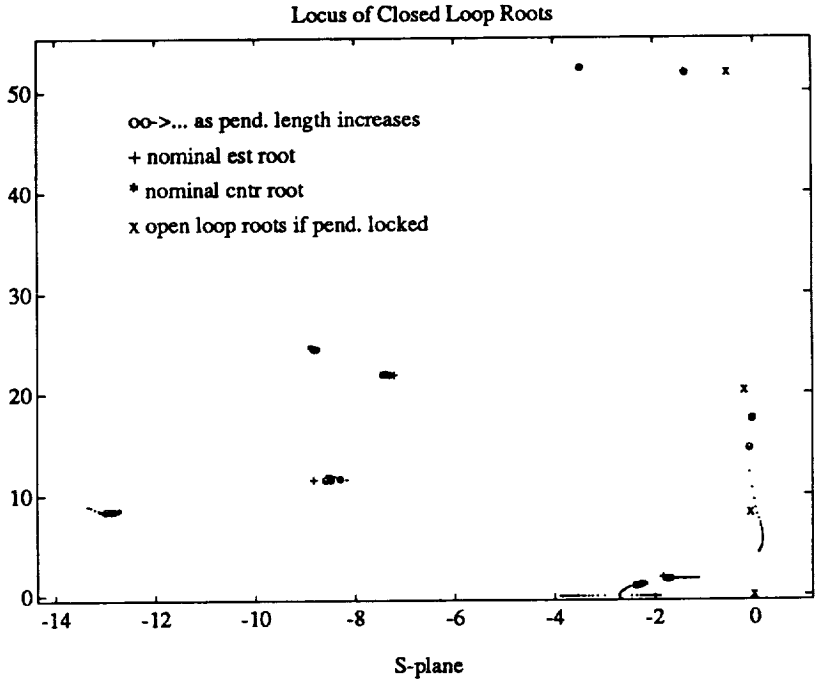

Figure 5.6: **LQG design w/o modelling pendulum** *For Experimental Hardware*

pendulum is very short, it represents a unmodelled high frequency part of the system. The roots of the controller are close to the design point, and the extra root of the system, due to the payload, is on the $j\omega$ axis at a high frequency. As the pendulum length increases, the controller roots move away from the design point, and the root due to the pendulum moves toward the origin along with bowing out towards the right half plane. This shows that the LQR controller is acting to undamp (destabilize) the pendulum.

Figure 5.6 shows a plot of the closed loop roots of the system for the actual experimental hardware as the pendulum length is varied. The nominal controller is an LQG (linear quadratic quassian) controller without modelling the internal dynamics of the payload. For some lengths of the pendulum, the controller adds damping to the pendulum and for other lengths the controller reduces the damping potentially driving it unstable. The pendulum mode appears to wiggle into the right half plane when the mode is near an arm mode and into the left half plane in between arm modes.

The conclusion is that if the control algorithm is designed for a rigid payload and the payload is not rigid, the system may in fact be unstable.

5.4 Experimental Results

This section of the report describes the experimental progress to date. The goal as stated in the introduction is to design a high-performance non-colocated controller that is capable of damping internal oscillations of the payload. For comparison purposes, Figure 5.7 shows the natural damping of the payload. The natural time constant of the pendulum is about 30 seconds.

The first controller is one where the dynamics of the payload are ignored. This control uses feedback of hub position, and endbody angle. The compensator is designed using LQG. The model of the system assumes that the pendulum is locked. Figure 5.8 shows the response of the pendulum and center of endbody to an initial condition. The damping time constant of the pendulum has been reduced in half to about 15 seconds. The pendulum frequency for this run is about 2.3 Hertz which is in between the first arm mode (1.3 Hz) and the second arm mode (3.2 Hz). Referring to Figure 5.6, the pendulum natural frequency is in a region where the controller should add slightly to the natural damping. Hence for this particular experimental configuration the payload damping ratio is increased by a controller ignoring the payload dynamics. However, for other frequencies of the pendulum the damping ratio will decrease.

To prove experimentally that it is possible to damp the payload reasonably, a non-colocated controller was designed using hub position, endbody angle, and pendulum position as feedbacks. Certainly, using pendulum position as a feedback is impractical for many situations, but it will allow a quick experiment to prove that damping is achievable. Again an LQG design is used. Figure 5.9 shows the response. The damping time constant of the pendulum has been cut to on the order of 2 seconds. This is more than a factor of 10 improvement on the natural damping.

Finally, a controller is designed where the pendulum angle is not measured directly.

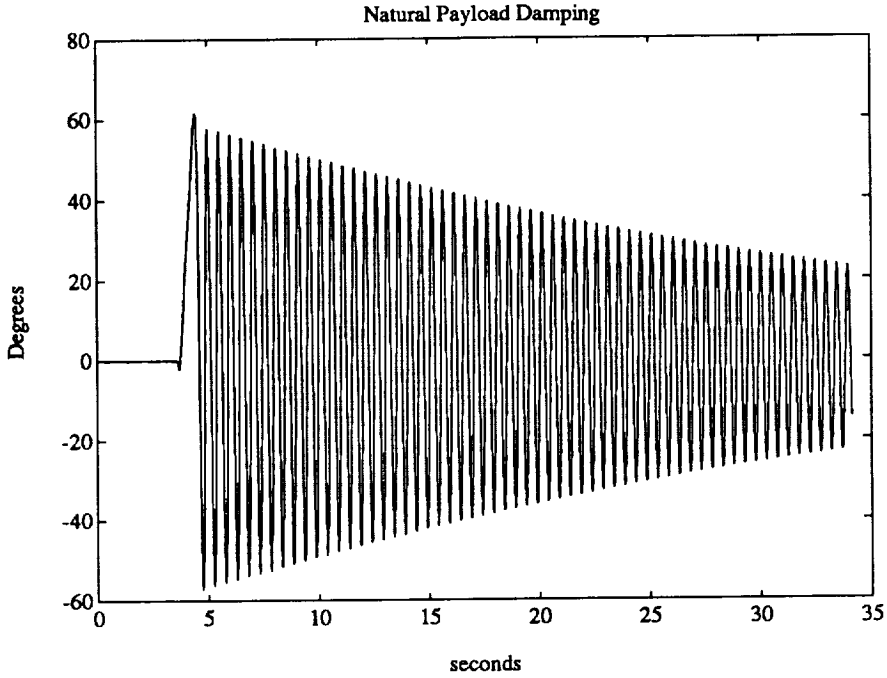

Figure 5.7: **Free Motion of Pendulum**

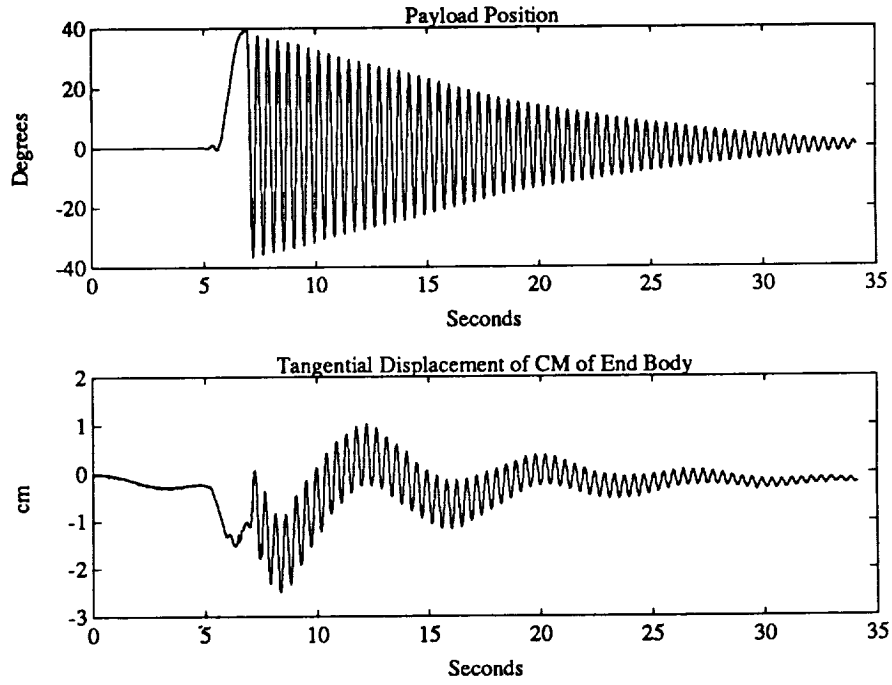

Figure 5.8: **Response With Dynamic Payload**

non-collocated control without accouting for payload

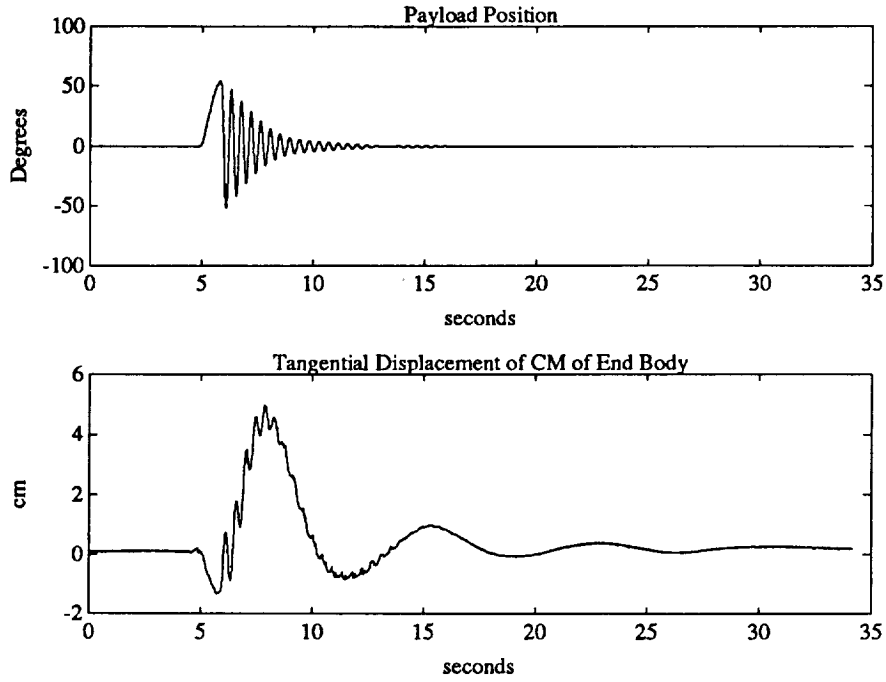

Figure 5.9: **Response With Dynamic Payload**

non-collocated control accounting for payload

The inputs to the compensator are the hub position, endbody angle, and the tangential displacement of the center of the endbody. Figure 5.10 shows the response. The time constant is still on the order of 2 seconds. However, there is a little bit of build up in the pendulum angle after it is initially damped. This is not fully understood yet. The ability to damp the pendulum without direct measurement is a major step forward.

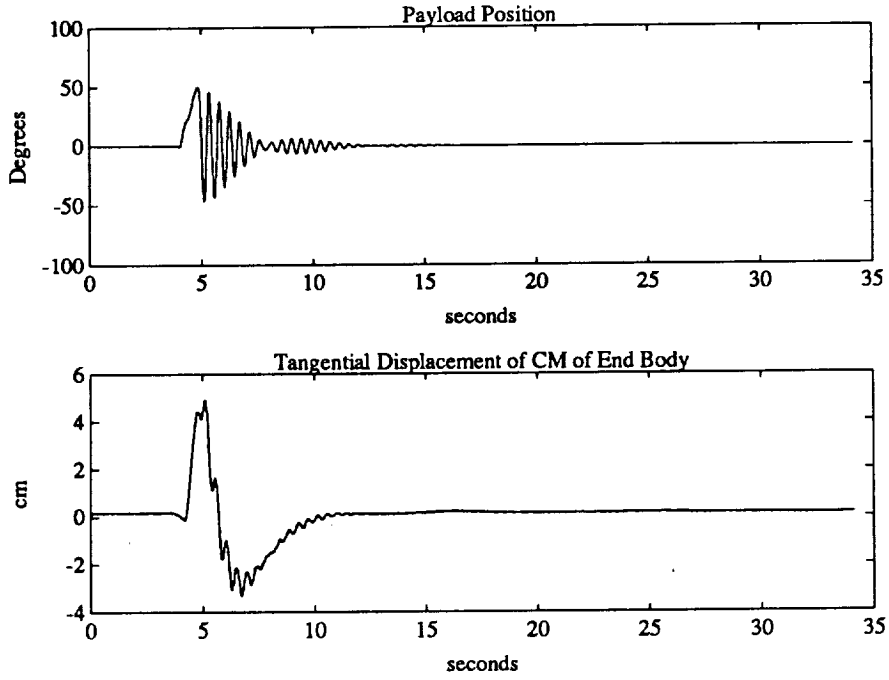

Figure 5.10: **Response With Dynamic Payload**

non-collocated control w/o measuring pendulum

5.5 Conclusions and Future Work

During the last six months, the experimental hardware for the dynamic payload experiment has been completed. The physically system has been mathematically modelled. The accuracy of the model has been verified by experimental data. Using the model, insight has been gained into the coupling between the dynamics of the payload and the arm. The model has also helped in sensor selection.

A stability analysis of the system has been performed. The analysis shows that for the general class of problems of flexible systems with unknown external (outside the sensor) dynamics controllers must account for the external dynamics or instability may result.

Finally, a non-colocated controller has been implemented that damps the pendulum a factor of 10 times faster than the natural damping without directly measuring the pendulum. Future work includes damping the pendulum without explicit knowledge of the pendulum frequency.

~

Bibliography

[1] Robert H. Cannon, Jr., Marc Ullman, William Dickson, Lawrence Alder, Edward Wilson, and Kurt Zimmerman. NASA Semi-Annual Report on Control of Free-Flying Space Robot Manipulator Systems. Semi-Annual Report 10, Stanford University Aerospace Robotics Laboratory, Stanford, CA 94305, February 1991.

[2] Robert H. Cannon, Jr., Marc Ullman, Ross Koningstein, William Dickson, Warren Jasper, and Lawrence Alder. NASA Semi-Annual Report on Control of Free-Flying Space Robot Manipulator Systems. Semi-Annual Report 9, Stanford University Aerospace Robotics Laboratory, Stanford, CA 94305, August 1990.

[3] Robert H. Cannon, Jr., Marc Ullman, Ross Koningstein, Stan Schneider, Warren Jasper, Roberto Zanutta, and William Dickson. NASA Semi-Annual Report on Control of Free-Flying Space Robot Manipulator Systems. Semi-Annual Report 7, Stanford University Aerospace Robotics Laboratory, Stanford, CA 94305, August 1988.

[4] Warren Jasper. *Experiments in Thrusterless Robot Locomotion Control for Space Applications.* PhD thesis, Stanford University, Stanford, CA 94305, September 1990. Also published as SUDAAR 595.

[5] Ross Koningstein. *Experiments in Cooperative-Arm Object Manipulation with a Two-Armed Free-Flying Robot.* PhD thesis, Stanford University, Department of Aeronautics and Astronautics, Stanford, CA 94305, October 1990. Also published as SUDAAR 597.

~

Appendix A

Experiments in Autonomonous Navigation and Control of a Mulit-Manipulator Free-Flying Space Robot

To appear as a chapter in
Space Robotics: Dynamics and Control
A collection of papers edited by Prof. Takeo Kanade and Dr. Yangsheng Xu, CMU
Published by Kluwer Academic Publishers

Experiments in Autonomous Navigation and Control of a Multi-Manipulator, Free-Flying Space Robot

Marc A. Ullman and Robert H. Cannon, Jr.
Aerospace Robotics Laboratory
Stanford University
Dept. of Aeronautics and Austronautics
Stanford, California 94305

Abstract

This paper reviews work performed at the Stanford University Aerospace Robotics Laboratory (ARL) in developing and controlling a multi-manipulator, free-flying space robot. The objective of this project was to create a laboratory version of a space robot capable of performing target tracking, acquisition, and manipulation. In particular, this paper focuses on the problems associated with capturing a free-floating object that is initially out of reach of the robot. A set of rules for generating an appropriate intercept trajectory is presented along with a controller architecture capable of carrying out the required actions. We conclude with a description of the physical hardware on which this approach was tested along with experimental data showing the successful capture of a spinning object.

Introduction

Although space presents us with an exciting new frontier for science and manufacturing, it has proven to be a costly and dangerous place for humans. It is therefore an ideal environment for sophisticated robots capable of performing tasks that currently require the active participation of astronauts.

As our presence in space expands, we will increasingly need robots that are capable of handling a variety of tasks ranging from routine inspection and maintenance to unforeseen servicing and repair work. Such tasks could be carried out by free-flying space robots equipped with sets of dextrous manipulators. These robots will need to be able to navigate to remote job sites, rendezvous with free-flying objects, perform servicing or assembly operations, and return to their base of operations. NASA's Manned Maneuvering Unit (MMU) was

built to outfit astronauts with similar capabilities. It enables them to carry out some of these tasks, but at a higher cost and with much greater risk to human safety than if robots were to be used.

Research Objectives

In order to advance the underlying theory and technology necessary for the aforementioned robotic capabilities to be realized, we have studied the problems associated with building and controlling autonomous free-flying space robots. Our objective was to demonstrate the ability to carry out complex tasks including acquisition, manipulation, and assembly of free-floating objects based on task-level commands. Our approach was to extend earlier ARL work in cooperative manipulation involving the use of fixed-base manipulators[1] to accommodate an actively mobile base thereby removing the workspace limitations inherent in the fixed-base implementation.

Experimental Verification

In order to test our design methodologies and control strategies, we have developed an experimental two-armed satellite robot that uses an air cushion support system to achieve—in two dimensions—the drag-free, zero-g characteristics of space. (See Figure 1.) The robot is a fully self-contained spacecraft possessing an on board gas supply for flotation and propulsion, rechargeable batteries for power, and on-board computers with sensing and driver electronics for navigation and control. Although the robot can function autonomously, its computers can also communicate with a network of workstations via a fiber optic Ethernet link.[1] An on-board camera provides optical endpoint and target sensing while an overhead global vision system facilitates robot navigation and target tracking.[2] We have used this system to demonstrate the successful intercept and capture of a free-flying, spinning object.

Background

A number of researchers have worked on the problem of controlling the endpoint position of a manipulator mounted on an *uncontrolled* free-flying base[2] [3][4][5][6][7]. Several have shown that with judicious path planning, the orientation of the base body can also be controlled. These approaches rely on the assumption that no external forces or torques act on the system thus leading to formulations based on conservation of the total linear and angular momentum. This assumption requires that the desired manipulator endpoint positions initially lie within reach and that the system starts out with zero linear and angular momentum.[3] In order to be truly useful, space robots will need to be

[1] It is our hope to replace this link with one of the new wireless LANs that are now coming to market, thereby making our robot truly autonomous.

[2] The global vision system serves as a convenient laboratory surrogate for a tracking system such as GPS that could be used for this purpose in space.

[3] These approaches could be extended to handle the case of a system having initial momenta such that, by the intercept time, the robot has coasted to within reach of the target; however, failing to explain how this set of circumstances comes about makes this an incomplete solution.

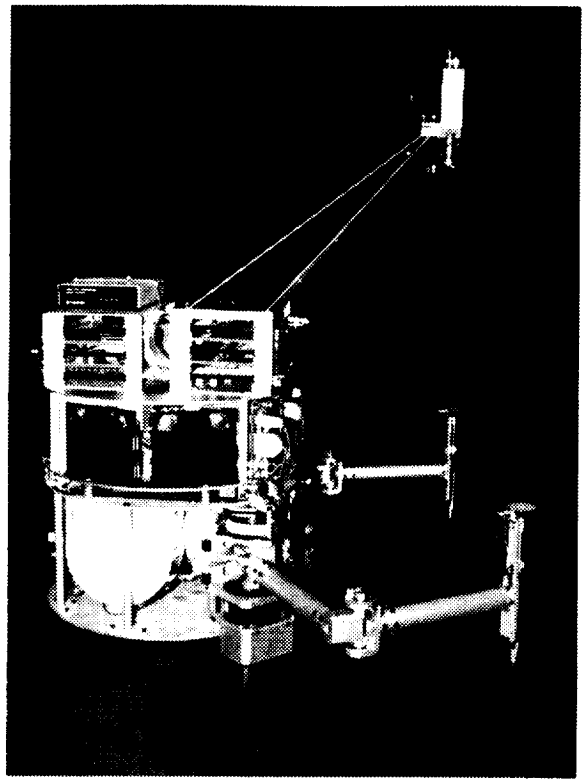

Figure 1: **Stanford Multi-Manipulator Free-Flying Space Robot**
This is a fully self-contained 2-D model of a free-flying space robot complete with on board gas, thrusters, electrical power, computers, camera, and manipulators. It exhibits nearly frictionless motion as it floats above a granite surface plate on a 0.005in thick cushion of air.

able to function in a much larger workspace than that of their immediate reach.

The execution of useful work in space requires the ability to simultaneously control both robot base and manipulator motions. Dubowsky, et al.[8] have recognized this requirement; however, they use gas jets for disturbance rejection rather than for actively controlling the base body to follow a specified trajectory. In general, rendezvousing with and capturing a free-flying object requires controlling both manipulator and base body positions to follow *coordinated* intercept trajectories.

Global navigation and control (or "gross motion" control) of a space robot therefore poses a set of interesting and unique challenges. These differ fundamentally from both the typical satellite positioning/attitude control problem and the case of a free-floating space robot with an uncontrolled base. This paper examines the problem of controlling the entire space robot system so as

to be able rendezvous and dock with moving objects.

Catching a Moving Object

In order to gain some additional insight into how to catch a moving object, we begin by looking at how humans address this task. When formulating a strategy for catching a moving object, one instinctively takes into consideration the object's mass and velocity, as the following examples show:

Example 1 From the game of baseball consider the case of a short stop fielding a grounder. Here the object being caught (the ball) has insignificant mass and inertia when compared to the player doing the catching. The short stop tries to execute the shortest, smoothest motion he can that allows him to get to the ball in time to pick it up and throw out the batter. He minimizes his body motion by reaching with his arms, typically following the minimum distance path to intersect the line of motion of the ball—one that is usually *perpendicular* to the ball's motion. That is, he lets the ball come to him and does not concern himself with trying to match the ball's velocity.

Example 2 By way of contrast, consider a person climbing on board a slowly moving train. In this case the person will have no noticeable effect on the motion of the train. He must match the velocity of the train by running along side it before grabbing on or he will likely suffer bodily injury. His intercept trajectory is therefore *parallel* to that of the object he is trying to catch.

In both cases the human must get his or her hands and body into proper position in order to effect the catch. As the mass of the object increases, so too does the necessity of matching the velocity of one's body with that of the object. Hence, one *coordinates* the motion of his or her arms and body.

Similarly, in order to successfully capture a free-floating target, a free-flying space robot must simultaneously control both its manipulators and its base body position and orientation. Since the corresponding states are coupled with each other it is necessary to view the system as whole rather than as two decoupled problems. Simply generating an realizable intercept trajectory—given the limited actuator authority available, the ever present dynamic constraints imposed by a free-floating robot, and any temporal constraints (e.g. the object might float beyond reach if not caught soon enough)—presents a formidable problem.

Intercept Trajectories

In order to rendezvous with and capture a moving target, a realizable intercept trajectory must be formulated. As was illustrated by our examples, the nature of this trajectory will vary according to the task at hand.

tip positions	Must always match with the object.
tip velocities	Must match with the object if it has any appreciable mass.
base position	Must be such as to insure that the object is within reach at intercept time.
base velocity	Must match that of the object if its mass is significant in comparison to the robot's.

Table 1: Intercept Trajectory Requirements

Trajectory Requirements

The intercept trajectory must always assure that the base position and orientation allow the manipulators sufficient freedom to successfully grapple the target without running into the limits of their workspace.

For low mass objects, the robot base intercept trajectory can simply be a straight line toward the object intercept point.[4] The grasped object will have very little effect on the robot motion and the arms will be able to position the object so as to keep it from contacting the robot base. For a massive object; however, the robot must carefully pull along side the object before attempting to grasp it to avoid the possibility of a collision. Thus, knowing the mass of the target allows us to optimize the necessary base motion. These ideas are summarized in Table 1.

Rules for Determining a Trajectory

By choosing the intercept time t_f we can describe the intercept requirements in terms of a set of rules. Knowing the intercept time allows us to predict the terminal object position and orientation (assuming that the object is not experiencing any unknown external forces and/or torques). We can then determine the peak acceleration requirements to see if they exceed our maximum actuator capabilities. If this is the case, the intercept time can be modified as necessary until an achievable path is obtained. These rules can be summarized as follows:

Rule 1 The required manipulator end point positions at intercept time are defined by the target position and orientation.

Rule 2 The manipulators should be in the center of their work space—this configuration allows the maximum range of motion for final correction of alignment errors.

The first two requirements constrain the desired base position at t_f to lie on a sphere (or circle in a 2-D world) whose radius R_d is defined as the distance from the center of the base to the manipulator tips. They also tell us the base orientation once its position is known, thus we have:

[4] Assuming that there are no obstacles in its path. Current work is addressing this issue.

Rule 3a If the object to be caught has insignificant mass, then we select the desired base position to be the point on the sphere closest to our initial base position (corresponding to the minimum distance path).

Rule 3b If the object to be caught is extremely massive, we consider the great circle defined by the intersection of the plane normal to the object's velocity vector at intercept time and the aforementioned sphere. The desired base position is selected to be the point on this great circle that is closest to our initial base position.

A non-linear weighting can be used to span the two extremes described by rules 3a and 3b. Similarly:

Rule 4 If the object to be caught has insignificant mass, we need not place any restrictions on the base velocity at intercept time. However, as the mass of the object increases, so too does the need to match the base velocity with that of the object.

Base Trajectories

The typical mode of transportation will be to use gas jet thrusters.[5] Since the robot base essentially behaves as a $1/s^2$ (double integrator) plant, the most fuel efficient trajectory is bang-off-bang for each of its degrees of freedom.[6] Clearly, the longer the robot can coast at peak velocity, the less fuel it needs to traverse a given distance in a fixed time. Thus we use a generic trajectory consisting of a linear function with parabolic blends described by the following family of equations:

$$a(t) = \begin{cases} a_1, & 0 \le t_1 \\ 0, & t_1 \le t \le t_2 \\ a_2, & t_2 \le t \le t_f \end{cases}$$

$$v(t) = \begin{cases} v_o + a_1 t, & 0 \le t_1 \\ v_o + a_1 t_1, & t_1 \le t \le t_2 \\ v_o + a_1 t_1 + a_2(t - t_2), & t_2 \le t \le t_f \end{cases}$$

$$s(t) = \begin{cases} s_o + v_o t + a_1 t^2/2, & 0 \le t_1 \\ s_o + v_o t - a_1 t_1^2/2 + a_1 t_1 t, & t_1 \le t \le t_2 \\ s_o + v_o t - a_1 t_1^2/2 + a_1 t_1 t & \\ \quad + a_2(t - t_2)^2/2, & t_2 \le t \le t_f \end{cases}$$

[5] In certain circumstances one can achieve motion through the use of the robot manipulator arms[9][10]

[6] For the 2-D case this concept is directly applicable since the three degrees of freedom $[x, y, \theta]$ and their derivatives can be specified independently. In the full 3-D case, this is no longer true since the final orientation is a function not only of the change in orientation angles but also of their respective time histories (i.e. it is a nonholonomic system). One partial solution to this problem is to find the *simple* rotation that takes the base from the given initial orientation to the desired final orientation. The final configuration would then be independent of the rate of rotation about this axis and we could optimize this trajectory. The drawback is that it makes no provisions for non-zero initial and/or final angular velocities.

The unknowns t_1 and t_2 along with the signs of a_1 and a_2 can be found to accommodate a specified set of initial and final conditions $[s_o, v_o]$ and $[s_f, v_f]$ respectively. By assuming that the magnitudes of a_1 and a_2 are known a priori,[7] the solution for the off- and on- times is given by a quadratic equation and the appropriate signs on the $a's$ are found via feasibility tests. Figure 2 shows an example of one such trajectory. We can also determine the minimum time (i.e. band-bang) path for a given maximum acceleration.

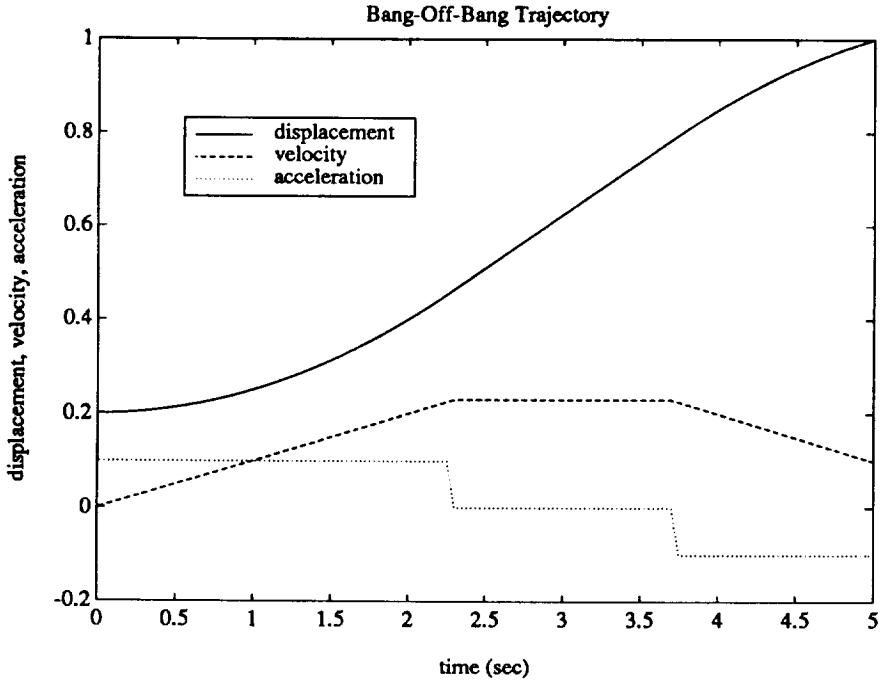

Figure 2: **Bang-off-bang Trajectory**

This example of a typical bang-off-bang trajectory time history for a single degree of freedom shows that we can satisfy arbitrary initial and final position and velocity conditions.

Manipulator Trajectories

When executing maneuvers in which the base moves a substantial distance, it is often desirable to control the motions of the manipulator endpoints in the reference frame of the robot base. This allows such actions as tucking the arms in to avoid bumping into other objects or holding them in a set position to steady a payload. When the robot gets "close" to a work site,

[7] We assume that a_1 and a_2 are equal in magnitude, however, the equations are general and do not require this assumption.

it can switch to an endpoint controller utilizing "operational" or workspace coordinates. These controllers require endpoint trajectories for the manipulator tips. Quintic polynomials trajectories are used since they allow the position, velocity, and acceleration of the manipulator endpoint to match the initial and final conditions at times of interest (e.g. mode switches, trajectory updates, and target intercepts).

Controller Architecture

Our controller architecture consists of a three-level hierarchy composed of a stateless remote graphical user interface, a high-level strategic controller, and a low-level dynamic controller based on an operational space computed torque formulation.

The graphical user interface (See figure 3) runs on a Sun Workstation and allows an operator to send high level commands to the robot by selecting icons and clicking on buttons.

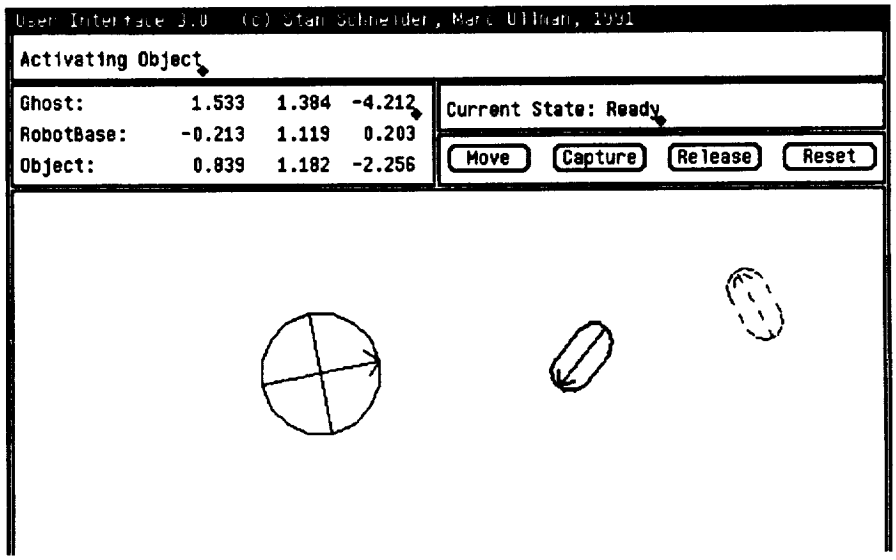

Figure 3: **Typical view of the user interface**

The graphical user interface provides "point and click" operation of the robot and allows the operator to control and monitor all operations remotely. Here a capture and move operation is underway.

The high-level strategic controller is based on a sophisticated finite state machine. It accepts commands from the remote user interface and reconfigures the low-level dynamic controller to carry out desired actions. A thorough discussion of the strategic controller is beyond the scope of this paper and can be found in [11].

The low-level dynamic controller uses an inverse dynamics or computed

torque formulation as described in the following sections.

Equations of Motion

Using Kane's method[12], the joint space equations of motion for the complete system can be expressed in terms of a vector of generalized coordinates $\mathbf{q}$ and a vector of generalized speeds $\mathbf{u}$. They are of the form:

$$\mathbf{F} = M(\mathbf{q})\dot{\mathbf{u}} + V(\mathbf{q}, \mathbf{u})$$

where $M(\mathbf{q})$ is the configuration dependent mass matrix, $V(\mathbf{q}, \mathbf{u})$ is the configuration and velocity dependent vector of non-linear terms, and $\mathbf{F}$ is the vector of generalized active forces. The generalized speeds are defined in terms of the state derivatives, $\dot{\mathbf{q}}$, by the relation

$$\mathbf{u} \triangleq Y(\mathbf{q})\dot{\mathbf{q}}$$

The generalized active forces are composed of linear combinations of the actuator forces and torques, τ, along with any external forces and torques (e.g. gravity), here assumed to be zero.[8]

$$\mathbf{F} = R(\mathbf{q})\tau$$

For our system consisting of the robot base and two manipulators, $\mathbf{q}$ is composed of the base position and orientation along with the joint angles describing the configuration of each manipulator arm.

$$\mathbf{q} = \begin{bmatrix} q_{base\_pos} \\ q_{base\_orient} \\ q_{right\_arm} \\ q_{left\_arm} \end{bmatrix}$$

Operational Space Computed Torque Controller

An operational space[13] computed torque controller facilitates specifying our desired behavior in Cartesian coordinates rather than in terms of the joint angles. This later approach is much preferred when using trajectories generated by path planning algorithms as described above.

The extension of the computed torque controller into operational (or Cartesian) space is fairly straight forward. We begin by defining a state vector of coordinates, $\mathbf{x}$, that describe the robot configuration in terms of variables we are directly interested in controlling. In our case, we have selected the base position and orientation and the manipulator endpoint positions so that:

$$\mathbf{x} = \begin{bmatrix} x_{base\_pos} \\ x_{base\_orient} \\ x_{right\_endpoint\_pos} \\ x_{left\_endpoint\_pos} \end{bmatrix}$$

[8] They are related by the partial velocities and partial angular velocities.

A set of kinematic equations relates this state vector to our original set of generalized coordinates $\mathbf{q}$. Typically one defines the relation between the time derivatives of these two state vectors (in an inertial reference frame) as the Jacobian yielding:

$$\dot{\mathbf{x}} = J\dot{\mathbf{q}}$$

However, since our equations have been cast in terms of generalized speeds, $\mathbf{u}$, we find it more convenient to make the following definition:

$$\mathcal{J} \triangleq JY^{-1}$$

so that

$$\dot{\mathbf{x}} = JY^{-1}\mathbf{u} = \mathcal{J}\mathbf{u}$$

Differentiating this relationship leads to

$$\ddot{\mathbf{x}} = \dot{\mathcal{J}}\mathbf{u} + \mathcal{J}\dot{\mathbf{u}}$$

from which we can solve for $\dot{\mathbf{u}}$

$$\dot{\mathbf{u}} = \mathcal{J}^{-1}(-\dot{\mathcal{J}}\mathbf{u} + \ddot{\mathbf{x}})$$

We can replace $\ddot{\mathbf{x}}$ with a desired acceleration $\mathbf{a}_{des}$ composed of both feedforward and feedback terms resulting from our commanded trajectory and our feedback control law respectively. With a proportional-derivative (PD) control law the desired acceleration vector is:

$$\mathbf{a}_{des} = \ddot{\mathbf{x}}_{cmd} + K_v(\dot{\mathbf{x}}_{cmd} - \dot{\mathbf{x}}) + K_p(\mathbf{x}_{cmd} - \mathbf{x})$$

where K_p and K_v are diagonal matrices containing the proportional and derivative feedback gains respectively.

Substituting the resulting expression for $\dot{\mathbf{u}}$ back into our original equations of motion yields a set of generalized forces from which we can determine the required actuator forces and torques.

$$\mathbf{F} = M(\mathbf{q})\{\mathcal{J}^{-1}(-\dot{\mathcal{J}}\mathbf{u} + \mathbf{a}_{des})\} + V(\mathbf{q}, \mathbf{u})$$

$$\tau = R^{-1}(\mathbf{q})\mathbf{F}$$

The actual implementation of the control algorithm described above involves a recursive formulation and is described in depth in [14].

On-Off Actuators

Although proportional thrusters are now available for critical applications we have assumed that the thrusters are of the on-off type capable of delivering only one level of thrust. Furthermore, we do not assume that the thrusters have been given linear characteristics though the use of an external pulse width modulation (PWM) mechanism. This assumption is based on the belief that the net impulse delivered during one sample period is of sufficiently fine granularity that we naturally get an inherently PWM-like behavior from the system.[9] We use a least squares table lookup to determine which thrusters to fire at each time step based on the controller force/torque output.

Experimental Work

In order to verify the utility of the control methodology described above, we have built a laboratory model of a two-armed space robot (see Figure 1) that experiences in two-dimensions the drag-free, zero-g characteristics of space. These characteristics are achieved using air cushion technology. The robot "floats" on a $9' x 12'$ granite surface plate with a drag-to-weight ratio of about 10^{-4} and gravity induced accelerations below $10^{-5}g$—a very good approximation to the actual conditions of space. In addition, this fully self-contained spacecraft possesses:

- an on board gas subsystem used both for flotation and for propulsion via thrusters,
- a complete electrical power system with plug-in rechargeable batteries packs[10] and power conditioning, distribution, and monitoring circuitry,
- a full complement of sensors and signal conditioning electronics,
- a high speed multi-processor computer system,
- a complete set of digital and analog data acquisition and control interfaces,
- a fiber-optic-based data/communications link to a network of off-board computers, and
- a camera and vision system for real-time tracking of the manipulator endpoints and target motions.

The robot measures $50cm$ in diameter and stands $65cm$ high with a total mass of just under $50kg$. In order to simplify maintenance operations as well as to facilitate future design modifications the robot was designed as a series of

[9] Since computed torque controllers are based on continuous time theory, they typically require high sample rates to insure the satisfaction of this assumption.

[10] The battery packs can also be recharged while on board the robot through the use of an umbilical power cord

independent modules. These modules take the form of layers, each of which performs a distinct function. The layers can be easily separated[11] when necessary for servicing or repair.

The base body has three degrees of freedom (x, y, θ) and sports eight gas jet thrusters mounted as four ninety-degree pairs sitting at the corners of a square inscribing its outer circumference.

A pair of two-link planer arms aligned with a set of ninety-degree separated rays are attached to the base. These manipulator arms are driven by a coaxial set of limited angle DC torque motors—the shoulder joint being driven directly while the elbow joint is driven though a cable from the elbow motor, which rides on the shoulder link. Both joints are instrumented with RVDTs for sensing joint angles. Analog differentiators provide corresponding rate signals in hardware. The manipulators are equipped with pneumatically actuated grippers that possess a single degree of freedom along the z-axis. Objects can be grasped by lowering the gripper plungers into cup-like grasp fixtures mounted on the objects.

The on-board computer system runs the VxWorks real-time operating system. This operating system allows us to develop code on Sun Workstations that can be downloaded to the target processors via a fiber optic Ethernet link. Since the real time operating system contains a complete implementation of TCP/IP and NFS the target processors can access files and data on our host server. We have configured the system as a set of subnets so that we can communicate between on- and off- board processors without incurring delays due to traffic on our workstation LAN.

A CCD camera-based vision system allows us to track the position of the robot and an array of targets in real-time at 60Hz. Unique patterns of three infrared LEDs are used to identify different objects. One camera mounted on board the robot provides precise information in the immediate workspace of the manipulators while a pair of cameras mounted above the granite table provides global positioning information for the robot and the target objects. Data from these off board cameras is relayed to the robot over the fiber optic LAN.

Experimental Results

Figure 4 shows the robot grasping a small target object. The object also floats on an air bearing supplied by a battery-powered aquarium air pump. The object can be sent across the granite table in a random direction with a rotation rate as high as 20 revolutions per minute. The robot can be commanded to capture and retrieve the object via the graphical user interface described in the section *Controller Architecture*. Figure 5 shows the time history of one such rendezvous. Since the object is of comparatively low mass ($\sim 1kg$), the robot follows a straight line path to intercept it. Upon grasping the object by inserting its "peg-in-the-hole"-style grippers, the robot brings the object smoothly to rest (in the robot's reference frame). It can then stow the object and transport it to some new location where it can release it.

[11] The main layers can be separated without any tools.

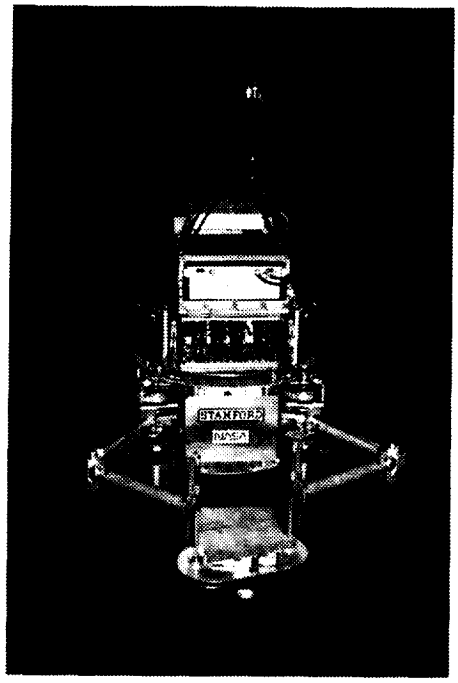

Figure 4: **Robot grasping floating object**
The robot can grasp the free-floating payload by inserting its plunger-type grippers into a pair of grasp fixtures.

Summary and Conclusions

This paper has described our work in developing a laboratory-based model of a multi-manipulator, free-flying space robot and the control strategies for enabling it to capture and manipulate free-floating targets. The problem of catching a moving object that is initially out of reach has been examined. It requires controlling the entire robot spacecraft as distinguished from the problem of controlling a set of manipulators on an uncontrolled free-floating base. In order to effect a successful catch, a realizable intercept trajectory must be found—one that takes into account the object's mass and inertia, as well as its velocity. We have presented a set of rules for determining such a trajectory along with a controller architecture capable of carrying out the necessary actions.

We have also described the design, construction, and testing of our 2-D model of a multi-manipulator, free-flying space robot. We have shown that by using the control approach presented in this paper, a fully autonomous rendezvous and capture operation can be performed, thereby providing a glimpse of the potential capabilities that free-flying space robots might someday possess.

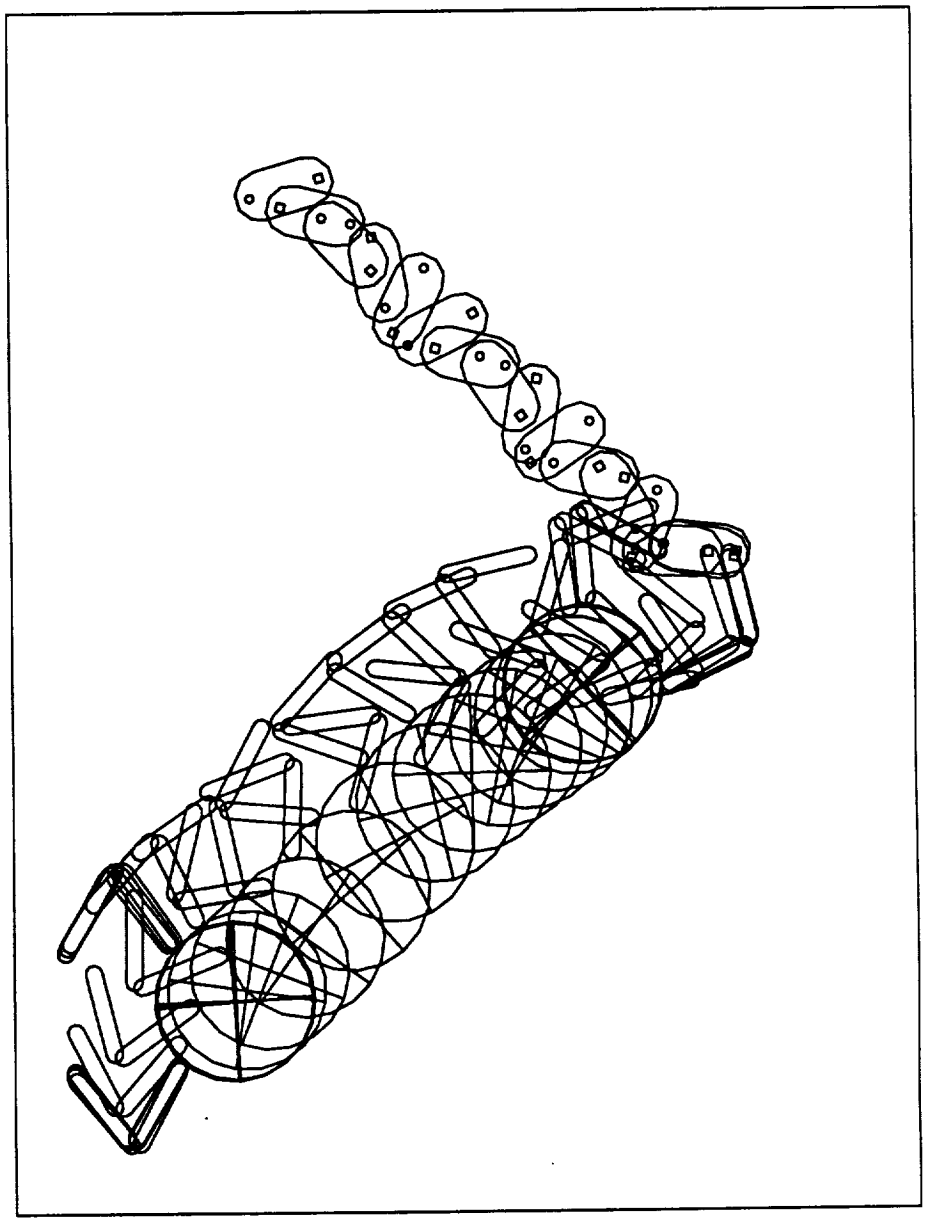

Figure 5: **Time history of object rendezvous and capture**
This "time-lapse" plot of experimental data shows the motion of a spinning object and the path the robot executed in order to intercept and capture it. The frame rate is 0.5Hz and this figure shows about 30 seconds of elapsed time.

Acknowledgments

This work was funded under NASA Cooperative Contract NCC 2-333.

References

[1] S. Schneider. *Experiments in the Dynamic and Strategic Control of Cooperating Manipulators.* PhD thesis, Stanford University, Stanford, CA 94305, September 1989. Also published as SUDAAR 586.

[2] Harold L. Alexander. *Experiments in Control of Satellite Manipulators.* PhD thesis, Stanford University, Department of Electrical Engineering, Stanford, CA 94305, December 1987.

[3] Ross Koningstein, Marc Ullman, and Robert H. Cannon, Jr. Computed torque control of a free-flying cooperating-arm robot. In *Proceedings of the NASA Conference on Space Telerobotics*, Pasadena, CA, February 1989.

[4] Yoji Umetani and Kazuya Yoshida. Experimental study on two dimensional free-flying robot satellite model. In *Proceedings of the NASA Conference on Space Telerobotics.* NASA Jet Propulsion Laboratory, 1989.

[5] Fumio Miyazaki Yasuhiro Masutani and Suguru Arimoto. Modeling and sensory feedback control for space manipulators. In *Proceedings of the NASA Conference on Space Telerobotics.* NASA Jet Propulsion Laboratory, 1989.

[6] Yoshihiko Nakamura and Ranjan Mukherjee. Redundancy of space manipulator on free-flying vehicle and its nonholonomic path planning. In *Proceedings of the NASA Conference on Space Telerobotics.* NASA Jet Propulsion Laboratory, 1989.

[7] E. Papadopoulos and S. Dubowsky. On the dynamic signularities in the control of free-floating space manipulators. In *Proceedings of the ASME Winter Annual Meeting*, pages 45–52, San Francisco, CA, December 1989. ASME.

[8] E. E. Vance S. Dubowsky and M. A. Torres. The control of space manipulators subject to spacecraft attitude control saturation limits. In *Proceedings of the NASA Conference on Space Telerobotics.* NASA Jet Propulsion Laboratory, 1989.

[9] Z. Vafa and S. Dubowsky. On the dynamics of manipulators in space using the virtual manipulator approach. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 579–585, Raleigh, NC, April 1987. IEEE, IEEE Computer Society.

[10] Warren Jasper. *Experiments in Thrusterless Robot Locomotion Control for Space Applications.* PhD thesis, Stanford University, Stanford, CA 94305, September 1990. Also published as SUDAAR 595.

[11] M. A. Ullman. *Experiments in Autonomous Navigation and Control of Multi-Manipulator Free-Flying Space Robots.* PhD thesis, Stanford University, Stanford, CA 94305, (July) 1991.

[12] Thomas R. Kane and David A. Levinson. *Dynamics: Theory and Application.* McGraw-Hill Series in Mechanical Engineering. McGraw-Hill, New York, NY, 1985.

[13] Oussama Khatib. Dynamic control of manipulators in operational space. In *Proceedings of the Sixth CISM-IFToMM Congress on Theory of Machines and Mechanisms*, pages 1128–1131, New Delhi, India, December 1983.

[14] Ross Koningstein. *Experiments in Cooperative-Arm Object Manipulation with a Two-Armed Free-Flying Robot.* PhD thesis, Stanford University, Department of Aeronautics and Astronautics, Stanford, CA 94305, October 1990. Also published as SUDAAR 597.